

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS ✓
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**

C# 920020049



**Europäisches
Patentamt**

**European
Patent Office**

**Office européen
des brevets**

Bescheinigung

Certificate

Attestation

Die angehefteten Unterla-
gen stimmen mit der
ursprünglich eingereichten
Fassung der auf dem näch-
sten Blatt bezeichneten
europäischen Patentanmel-
dung überein.

The attached documents
are exact copies of the
European patent application
described on the following
page, as originally filed.

Les documents fixés à
cette attestation sont
conformes à la version
initialement déposée de
la demande de brevet
européen spécifiée à la
page suivante.

Patentanmeldung Nr. Patent application No. Demande de brevet n°

03405168.0

Der Präsident des Europäischen Patentamts;
im Auftrag

For the President of the European Patent Office

Le Président de l'Office européen des brevets
p.o.

R C van Dijk



Europäisches
Patentamt

European
Patent Office

Office européen
des brevets

Anmeldung Nr:
Application no.: 03405168.0
Demande no:

Anmeldetag:
Date of filing: 12.03.03
Date de dépôt:

Anmelder/Applicant(s)/Demandeur(s):

International Business Machines Corporation
New Orchard Road
Armonk, NY 10504
ETATS-UNIS D'AMERIQUE

Bezeichnung der Erfindung/Title of the invention/Titre de l'invention:
(Falls die Bezeichnung der Erfindung nicht angegeben ist, siehe Beschreibung.
If no title is shown please refer to the description.
Si aucun titre n'est indiqué se référer à la description.)

Method for monitoring events in a computer network

In Anspruch genommene Priorität(en) / Priority(ies) claimed / Priorité(s)
revendiquée(s)

Staat/Tag/Aktenzeichen/State/Date/File no./Pays/Date/Numéro de dépôt:

Internationale Patentklassifikation/International Patent Classification/
Classification internationale des brevets:

G06F1/00

Am Anmeldetag benannte Vertragsstaaten/Contracting states designated at date of
filing/Etats contractants désignées lors du dépôt:

AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HU IE IT LU MC NL
PT SE SI SK TR LI

Method for monitoring events in a computer network

The present invention relates to a method for monitoring events in a computer network, said computer network triggering said events, wherein each event is provided with attribute values allocated to a given set of attributes.

With the expansion of the internet, electronic commerce and distributed computing, the amount of information transmitted via electronic networks is continuously increasing. Such possibilities have opened many new business horizons. However, they have also resulted in a considerable increase of illegal computer intrusions.

An emerging trend that addresses this problem is the development of intrusion detection systems. These systems are aimed to detect attacks on the computer network by monitoring all network activities. Network activities are usually monitored by the intrusion detection system as a time-ordered sequence of events wherein each event is characterized by a given set of attributes, so-called dimensions. Each event therefore forms an n-dimensional space.

The monitoring of a high number of events each having many attributes triggered by an intrusion-detection system is a task that requires high skill and attention from the monitoring staff, since a large fraction of the triggered events is regularly reported. The challenge for an operator of the intrusion detection system is to spot those events that are indicators of a real security problem. In order to distinguish security problem events from "false positive" alarms, the operators of the intrusion detection system usually watches out for interesting event patterns by means of a pattern detection algorithm. This pattern detection algorithm enables to detect whether an arrived event is part of a given pattern on the basis of a comparison of the attributes allocated to this given pattern and the attributes assigned to the arrived

event. For example, a pattern detection algorithm may determine whether the events triggered by the intrusion-detection systems all involve the same source IP, i.e. involve the same attacking machine, or the same destination IP, i.e. involve the same attack machine.

In order to render it possible for the operator to supervise the events triggered by the intrusion-detection system a suitable event visualization is needed. Current intrusion event presentation methods can be classified into three different groups: a first group of methods provides the operator of the intrusion detection system with a tabular text display of the relevant event information. For example, the operator console so-called Event Viewer of IBM Tivoli Enterprise Console TEC uses such a presentation method. In order to distinguish "false" positive events from real security problem events, a time-consuming comparison of textual information has to be carried out, making it difficult to spot interesting event patterns.

A second group of prior art event visualization methods provides the operator of the intrusion-detection system with a graphical representation of event information, but does not present the arrival time of the events. This second group method renders it possible to present various relations between event attributes. Such a second group method is known from Erbacher et al., Intrusion and Misuse Detection in Large-Scale Systems, IEEE CGA (2002). This document describes a visualization method representing security events as lines between points, each point representing a specific originating IP address or a specific destination IP address. From Girardin et al., A Visual Approach for Monitoring Logs, Proc. 12th Usenix System Administration Conference, Boston, Massachusetts, USA, 1998, a further second group method is known using a parallel coordinate visualization technique to represent different attributes of events. The disadvantage of the second group methods is that they do not display the event

time, which is the most important event attributes. This makes it difficult for operators of the intrusion-detection system to quickly orient themselves if they have not watched the display for a while.

5

A third group of prior art event monitoring methods enables an event visualization that represents the arrival time of events as a separate event attribute. The arrival time of the event is regularly displayed as the x-axis of cross-plot.

10 From Ma et al., Event Miner: An Integrated Mining Tool for Scalable Analysis of Event Data, May 2002, a visualization method is known using a two-dimensional mapping technique of arbitrary event attributes versa arrival time enabling an operator to analyze the event history. The disadvantage of this
15 method is that only one of the event attributes may be plotted versus the arrival time of the events. Thus, the operators have to switch continuously between the various event attributes to make sure that they do not miss a significant event pattern. From Haines et al., Visualization Techniques
20 for Event Stream Analysis, Eurographics UK Chapter 15th Annual Conference, Norwich, 1997, an event visualization technique is known using a vertical stack of cross plots to display multi-event attributes versus event arrival time. This known visualization technique works well if only a few event
25 attributes have to be monitored simultaneously on a screen. A problem may, however, occur if an operator of the intrusion detection system has to supervise a large number of event attributes. He then has to simultaneously watch a large number of different plots each displaying an event attribute versus
30 the event arrival time. In consequence, a high attention of the operator is required to detect all the security problems derivable from the displayed events.

In the view of the foregoing, an object of the present invention
35 is to provide a method of monitoring events in a computer network enabling an operator of an intrusion-detection

system to simultaneously monitor various event attributes versus the arrival time of the events.

This object is met by a method of monitoring events in a
5 puter network according to claim 1. Preferred embodiments disclosed in the dependent claims.

The inventive method of monitoring events in a computer network, said computer network triggering said events, each
10 event being provided with attribute values allocated to a given set of attributes includes the steps of providing an event display with a cross plot having two coordinate axes the x-axis presenting a time period and the y-axis presenting an attribute value range, determining a primary attribute
15 the events selected from the given set of attributes to be presented with its attribute values on the y-axis of the cross plot, allocating a first display label to the events indicating the attribute values of the primary attribute, providing a pattern algorithm to detect whether an arrived
20 event is part of a given pattern on the basis of a comparison of the attributes allocated to the given pattern and of the attributes assigned to the arrived event, providing a mapping algorithm to map any attribute value of an attribute selected from the given set of attributes onto the y-axis of the cross
25 plot, allocating a second display label to the events indicating the attribute value of the attributes being uncovered as part of the given pattern, plotting all the events arrived within the time period and including an attribute value allocated to a primary attribute into the cross plot with the
30 first display label indicating the primary attribute, the position of the first display label of each event in the cross plot being determined on the basis of the attribute value of the primary attribute of the event and its arrival time, and plotting all the events arrived within the time period and
35 being detected by the pattern algorithm as part of the given pattern into the cross plot with the second display label indicating the given pattern, the position of the second display

play label of each event in the cross plot being determined by the mapping algorithm on the basis of the attribute value of the attribute of the event as being uncovered as part of the given pattern and its arrival time.

5

The inventive event visualization method only renders it necessary for an operator of the intrusion-detection system to supervise one single cross plot, which displays all relevant events. The x-axis of the cross plot of the event display indicates the arrival times of the relevant events. The y-axis represents the primary attribute values of the events in which the examiner is mainly interested. Additionally, all the events being detected by the pattern algorithm as part of an interesting event pattern are displayed in the cross plot.

10

In order to differentiate the events associated with the primary attribute from the events being part of the interesting event pattern, a first display label is assigned to all events including a primary attribute value and a second display label is assigned to all events indicating the attribute values of the attributes being uncovered as part of the relevant event pattern. By using the inventive method of monitoring events, the event display presents a plot of information of the main event attribute versus the arrival time of the event by using a first display label for the plotted events wherein the interesting event pattern derived from other event attributes is simultaneously presented by using the second display label for these events. If the operator of the intrusion detection system wants to investigate the events being detected as part of a given pattern in more detail, he can easily switch to the corresponding event attribute by selecting a mark of the second display label in the cross plot.

15

20

25

30

According to a preferred embodiment, the attribute values and the arrival time of a new event are recorded, on the basis of the recorded attribute values of the event it is determined whether or not the newly arrived event includes an attribute value of the primary attribute and if the newly arrived event

35

includes such an attribute value, the x-axis of the cross plot is shifted so that the time period being presented on the x-axis covers the arrival time of the event so that the events arrived within the shifted time period may be plotted into the cross plot with the first display label indicating their primary attribute values. This performance enables a fast display of the events including the primary attributes.

According to a further preferred embodiment, it is determined on the basis of a recorded attribute value of a newly arrived event whether or not the newly arrived event is part of a given pattern on the basis of a comparison of the attributes allocated to a given pattern and of the attributes assigned to the arrived event. If the newly arrived event includes an attribute value of the given pattern, the newly arrived event is added to the previous events being detected as part of the given pattern and all the events being associated with the given pattern are re-drawn in the cross plot. This technique enables a fast display of the events associated with an interesting event pattern.

Moreover, if a newly arrived event does not include an attribute value of the given pattern it is preferred to determine on the basis of recorded attribute values of all previous arrived events by means of the pattern algorithm whether or not a newly arrived event is part of a new pattern on the basis of a comparison of the attributes allocated to the pattern and of the attributes assigned to the arrived event. If the newly arrived event forms a new pattern together with the previously recorded events, a third display label is located to the events indicating the attribute values of the attributes being uncovered as part of the new pattern. Then all the events being detected by means of the pattern algorithm as part of the new pattern are plotted into the cross plot with a third display label indicating the new pattern. This technique enables that the event display always presents

all event patterns in all attribute dimensions independent from the actually selected dimension.

Moreover, according to another preferred embodiment, if t
5 an operator wants to change the primary attribute to be d
played on the event display and therefore switches to ano
event attribute, all the events labels are removed from t
cross plot. Then a further display label is allocated to
events indicating the attribute values of the new primary
10 tribute. Finally all the events arrived within the time p
riod presented on the x-axis of the cross plot and includ
an attribute value of the new primary attribute are plott
into the cross plot with the further display label indica
the new primary attribute. This technique enables the ope
15 tor a fast change between interesting attributes of event
triggered by the computer network.

According to another preferred embodiment, if the operato
selects one of the events, e.g. by moving the cursor near
20 over the plotted event display label, all the attribute v
ues recorded for this event are plotted into the cross pl
with their respective display labels. Moreover, textual i
formation associated with the selected event may be displ
on the event display. This technique enables the operator
25 quickly obtain all the information necessary to evaluate
interesting event.

According to another preferred embodiment, the pattern al
rithm is suitable to perform multi-attribute pattern reco
30 tion so that various interesting event patterns may be si
taneously displayed in the cross plot. In order to improv
the visualization of the pattern, it is further preferred
that all the events uncovered as part of the pattern are
clustered by a corresponding display label to distinguish
35 interesting event pattern from other patterns. The present
tion of the events is further improved by using display la

bels for indicating the events in the cross plot including specific color and/or a specific mark layout.

5 The foregoing and other objects, features and aspects and advantages of the present invention will become more apparent from the following detailed description of the present invention when taken in conjunction with the accompanied drawings.

10 Figure 1 is a conceptual view on the inventive method of monitoring events in a computer network;

Figure 2 is an inventive processing flow to display a newly arrived event;

15 Figure 3 is a processing flow for a user input to switch primary attribute of the events to be displayed;

Figure 4 is a processing flow for a user input to select specific event to be displayed in detail; and

20 Figure 5 is a data-flow diagram disclosing the functional components involved in generating the inventive event visualization.

25 Carefully logging network activities is essential to meet requirements of high security and optimal resource availability. However, detecting break-in attempts within the network activities is a difficult task. Making the distinctions between misuse and normal use and identifying intrusions using novel attack techniques is difficult.

30 The invention deals with an improved visual approach for monitoring events triggered by one or more intrusion detection systems in a computer network. However, the inventive technique may also be useful for displaying other types of events, not just intrusion events.

The monitoring of events, in particular intrusion events, is a task that requires high skill and attention from the monitoring staff. The reason for this is that a large fraction of the reported events are simply so-called "false" positive alarms. The challenge for the operator is therefore to spot those events that are associated with a real security problem. In order to identify such security events, the operator of the intrusion detection system is on the one hand interested in continuously watching a main characteristic of the incoming events and on the other hand to uncover interesting event patterns. Intrusion detection systems normally generate events provided with attribute values allocated to a given set of attributes to supervise the network activities. These attributes are frequently called dimensions.

It is the object of present intrusion detection visualization technique to display event information in such a way that it makes easy for an operator to distinguish false positive events from events belonging to a security problem. The inventive visualization technique, which is detailed below performs a visual fusion of multi-event attributes on a single display. The inventive method improves the state of the art by helping the operator to become aware of all relevant event patterns while looking only at a single monitor screen without the need to cycle around through multiple displays.

According to the invention, events which are triggered in a computer network, each event being provided with values allocated to a given set of dimensions, are monitored with a cross plot having two coordinate axes, the x-axis presenting a time period and the y-axis presenting a selected dimension value range. The operator determines a primary dimension of the events selected from the given set of dimensions to be presented with its dimension values on the y-axis of the cross plot. This primary dimension is associated with a first unique label, preferably a unique color or a unique mark layout. Moreover, it is preferred that each dimension of the

given set of dimensions is associated with a unique label. Moreover, a pattern algorithm is provided in the event monitoring device to detect whether an arrived event is part of given pattern on the basis of a comparison of the dimensions allocated to the given event pattern and the dimensions assigned to an arrived event. It is preferred that the pattern algorithm is able to simultaneously detect a multitude of event patterns. Moreover, the event monitoring device is provided with a mapping algorithm to map any dimension value of a dimension selected from the given set of dimensions onto the dimension value range of the selected primary dimension presented on the y-axis of the cross plot.

The event visualization is performed in that all events arrived within the time period presented on the x-axis of the cross plot and including a dimension value allocated to the primary dimension are plotted into the cross plot with the corresponding display label indicating the primary dimension. The position of the display label of each plotted event is determined on the basis of the corresponding dimension value of the primary dimension of the event and its arrival time. Further, all the events that arrived within the time period presented on the x-axis and being detected by means of a pattern algorithm as part of the given pattern, are also plotted into the cross plot with a unique second display label indicating the given pattern. The second display label indicating the pattern is preferably an additional mark layout combining all the events corresponding to the pattern in the cross plot. The position of the second display label of pattern events in the cross plot is determined by the mapping algorithm on the basis of the dimension values of the dimensions of the events being uncovered as part of the pattern and their arrival time.

Figure 1 presents a series of eight events E_n to E_{n+8} being recorded one after the other by the inventive event visualization device. Each event is associated with a set of dimen-

sions p, three dimensions p1 to p3 being indicated. Figure 1 shows a time vector on which the arrival time of each event E_n to E_{n+8} is marked. Below the time vector, Figure 1 further shows three cross plots, the x-axis of each cross-plot presenting a time period and the y-axis of each cross-plot presenting a dimension value range for dimensions p1 to p3, respectively. In the first cross plot, all the events arrived within the time period and including a dimension value allocated to the dimension p1 are plotted with a first color. The same applies to all the events including a dimension value allocated to the dimension p2 in the second cross plot and to all the events including a dimension value allocated to the dimension p3 in the third cross plot.

In the embodiment presented in Figure 1, the operator has determined dimension p1 of the recorded events as the primary dimension. In consequence the pattern algorithm explores whether any of the dimensions p1 to p3, are covered by a given pattern. For example the pattern algorithm examines whether all the events involve the same source IP and the same destination IP. All the events uncovered as part of the given pattern are connected with lines, as shown in the second cross plot and the third cross plot.

All the three cross plots p1 to p3 are finally combined to one single cross plot shown at the bottom of Figure 1, wherein all the events arrived within the time period and including a dimension value allocated to the primary dimension p1 are plotted with the associated unique color and mark layout. Further, all the events arrived within the time period and being detected by the pattern algorithm as part of the given pattern, are plotted into the cross plot with their unique colors indicating the respective dimensions of the pattern wherein all the events of the pattern are connected with lines.

The inventive method of event visualization enables the operator with a single view onto the x-y-coordinate system to monitor all the relevant events occurring in a computer network. The inventive technique provides the possibility that
5 the operator may look at any time at a plot of information dealing with one primary event dimension. These events are plotted with a unique display label. Moreover, all the interesting event patterns of the other dimension plots superimpose this primary dimension plot indicated by their corresponding unique display labels.
10

Figure 2 presents a processing flow for a newly arrived event. If a new event E_n arrives (step S1), the dimension values and arrival time of the newly arrived event are recorded. Furthermore, on the basis of the recorded dimension
15 values, it is determined whether or not the newly arrived event includes a dimension value of the primary dimension. If the newly arrived event includes a dimension value of the primary dimension, in step 2 the event display is shifted to make room for the plot of the newly arrived event, i.e. the
20 x-axis of the event display is shifted so that the time period presented on the x-axis of the plot covers the arrival time of the newly arrived event. Moreover, all the events which are recorded before the new time period presented on the x-axis are removed. This also applies to all the patterns
25 without any current events within the time period presented on the x-axis of the cross plot. In the next step S3, the newly arrived event is plotted into the cross plot with the unique color associated with the primary dimension. Then in step 4, on the basis of the recorded dimension value of all
30 previously arrived events, it is determined by means of the pattern algorithm whether the newly arrived event is part of the given pattern on the basis of a comparison of the dimensions allocated to the given pattern and the dimensions assigned to the newly arrived event. If the newly arrived event
35 includes a dimension value of the given pattern, the event is added in step 5 to the previous events being detected as part

of the given pattern and all these events being associated with the given pattern are re-drawn in the cross plot.

5 If the newly arrived event does not include a dimension value of the given pattern, it is determined in step S6 on the basis of the recorded dimension values of the previously arrived events by means of the pattern algorithm whether or not the newly arrived event is part of a new pattern on the basis of a comparison of the dimensions allocated to the new pattern and the dimension values assigned to the arrived event. 10 If the newly arrived event forms a new pattern together with the previously recorded events, all the events detected as part of the new pattern are plotted into the cross plot with their unique colors corresponding to the respective dimensions (step S7). If no new pattern is detected, the program 15 flow is terminated (step S8).

Figure 3 shows a program flow enabling the operator to change the primary dimension to be displayed. In a first step S11, 20 the operator switches the primary dimension to be displayed. In the next step S12, the new primary dimension is selected. The program then clears the display (step S13) and plots all the events arrived within the time period and including a dimension value allocated to the new primary dimension into the 25 cross plot with a corresponding display label indicating the new primary dimension (step S14). Then, all the detected patterns are also plotted into the cross plot (step S15).

30 If the operator intends to investigate the context of the pattern in more detail, a program flow may take place as shown in Figure 4. The operator may move the cursor to a plotted dot in the display and selected this dot (step S21). In the next step S22, the program plots all the dimension information into the cross plot corresponding to the selected 35 event. Further, a full picture of the event is displayed in a further step S23 by presenting a textual representation of all the event properties. The textual representation of the

event properties can be provided either in a separate window or by labeling all the displayed event dots. The step S23 may be triggered separately by the operator, for example, with a further push of a mouse key, when the cursor controlled by the mouse is located at the plotted dot. It is possible that the operator may select multiple events, for example, by shift clicking.

Figure 5 shows a data flow diagram presenting the functional components involved in the inventive event visualization technique. The central device 1 is the event dimension/display mapping component. The central device 1 takes the following information as an input: Information on detected event patterns from a pattern detector 2. Further, mapping definition information as input from a corresponding mapping database 3. This information specifies a function for each event dimension that maps any event dimension value into a value range of the y-axis of the corresponding event display x-y-coordinate system. In order to carry out this mapping performance, the mapping definition information specifies a family of functions m with individual functions $m_{\text{dimension}}$: $\text{domain}_{\text{dimension}} \rightarrow Z$. Further, the central device 1 receives information on the current selected primary event dimension 4 to be displayed and information on the current event from the event database 5. Said event database 5 is also connected to the pattern detector 2. On the basis of the input information, the central device 1 determines the events and the patterns to be displayed and output the data to be displayed to the event and pattern display 6. Said event and pattern display 6 enables an interaction with the operator, the operator interaction may affect the event database 5 and/or the selected dimension 4.

Figure 1 of the present application shows as an example a linear pattern, i.e. all dots are located on a single row which is detected by the pattern algorithm and visualized. However, also more complex dimension patterns can be detected

by the pattern detection algorithm and be displayed in a similar manner, as shown in Figure 1. To present a complex pattern, the display technique may highlight the involved event dots and possibly connect them with a polygon line to emphasize the pattern. The inventive method not only performs "within dimension" patterns, but also may use an algorithm to detect multi-dimension patterns. The pattern detection algorithm might further use background information such as the operating system, vulnerabilities of the attacked machine as well as other information gathered from a network security scan. It is also possible to integrate such event background information as additional displayable event dimensions.

A problem with plotting information on multi-dimensions into a single cross plot may be that the dots can be clustered and occlude each other. To reduce such a clustering of the displayed dimensions, it may be possible to assign a unique y-position to each dimension.

Claims

1. A method of monitoring events in a computer network, said computer network triggering said events, each event being provided with attribute values allocated to a given set of attributes, comprising the steps of providing an event display with a cross plot having two coordinate axes, the x-axis presenting a time period and the y-axis presenting an attribute value range, determining a primary attribute of the events selected from the given set of attributes to be presented with its attribute values on the y-axis of the cross plot, allocating a first display label to the events indicating the attribute values of the primary attribute, providing a pattern algorithm to detect whether an arrived event is part of the given pattern on the basis of a comparison of the attributes allocated to the given pattern and of the attributes assigned to the arrived event, providing a mapping algorithm to map any attribute value of an attribute selected from the given set of attributes onto the y-axis of the cross plot, allocating a second display label to the events indicating the attribute values of the attributes being uncovered as part of the given pattern, plotting all the events arrived within the time period and including an attribute value allocated to the primary attribute into the cross plot with the first display label indicating the primary attribute, the position of the first display label of each event in the cross plot being determined on the basis of the attribute value of the primary attribute of the event and its arrival time, and plotting the all events arrived within the time period and being detected by means of the pattern algorithm as part of the given pattern into the cross plot with the

second display label indicating the given pattern, the position of the second display label of each event in the cross plot being determined by the mapping algorithm on the basis of the attribute value of the attribute of the event being uncovered as part of the given pattern and its arrival time.

2. The method according to claim 1 comprising the further steps of
recording the attribute values and the arrival time of a new event,
determining on the basis of the recorded attribute values of event whether or not the newly arrived event includes an attribute value of the primary attribute,
if the newly arrived event includes an attribute value for the primary attribute shifting the x-axis of the cross plot so that the time period being presented on the x-axis covers the arrival time of the event, and plotting the event arrived within the shifted time period into the cross plot with the first display label indicating the primary attribute.

3. The method according to claim 2 comprising the further steps of
determining on the basis of the recorded attribute values of event whether or not the newly arrived event is part of the given pattern on the basis of a comparison of the attributes allocated to the given pattern and of the attributes assigned to the arrived event,
if the newly arrived event includes an attribute value of the given pattern adding the event to the previous events being detected as part of the given pattern, and redrawing all the events being associated with given pattern in the cross plot.

35

4. The method according to claim 3 comprising the further steps of
if the newly arrived event does not include an attribute value of the given pattern,
5 determining on the basis of the recorded attribute values of all previous arrived events by means of the pattern algorithm whether or not the newly arrived event is part of a new pattern on the basis of a comparison of the attributes allocated to the new pattern and of the attributes assigned to the arrived events, if the newly
10 arrived event forms together with previous recorded events the new pattern, allocating a third display label to the events indicating the attribute values of the attributes being uncovered as part of the new pattern, and plotting the all events being detected by means of the
15 pattern algorithm as part of the new pattern into the cross plot with the third display label indicating the new pattern, the position of the third display label of each event in the cross plot being determined by the mapping algorithm on the basis of the attribute value of the attribute of the event being uncovered as part of the new pattern and its arrival time.

5. The method according to any of claim 1 to 4 comprising the further steps of
25 removing all the events including an attribute value allocated to the primary attribute from the cross plot, if a primary attribute to be presented with its attribute values on the y-axis of the cross plot is changed, allocating a fourth display label to the events indicating the attribute values of the new primary attribute,
30 and plotting all the events arrived within the time period and including an attribute value allocated to the new primary attribute into the cross plot with the fourth display label indicating the new primary attribute, the position of the fourth display label of each event in
35

the cross plot being determined on the basis of the attribute value of the primary attribute of the event and its arrival time.

- 5 6. The method according to any of claim 1 to 5 comprising
the further steps of
plotting all attribute values recorded for an event with
the respective display label into the cross plot if the
event is selected by an operator, and
10 displaying textual information associated with the selected event on the event display.
- 15 7. The method according to any of claim 1 to 6, wherein the
pattern algorithm is suitable to perform multi-attribute
pattern recognition.
- 20 8. The method according to any of claim 1 to 7, wherein
each display label includes a specific color and/or a
specific mark layout.
- 25 9. The method according to any of claim 1 to 8, wherein all
events being uncovered as part of the pattern are clustered by the corresponding display label.
- 30 10. A computer program containing a program code to carry
out the steps of the method of any of claims 1 to 9,
when the program code is running on a computer.
- 35 11. A computer program containing a program code to carry
out the steps of the method of any of claims 1 to 9,
said program code being stored on data carrier.
12. An event visualization device for monitoring events in a
computer network, the device comprising means to perform
the steps of the method as claimed in claims 1 to 9.

Abstract

The present invention relates to a method of monitoring events in a computer network, said computer network triggering said events, each event being provided with attribute values allocated to a given set of attributes, which includes the steps of providing an event display with a cross plot having two coordinate axes, the x-axis presenting a time period and the y-axis presenting an attribute value range, determining a primary attribute and a corresponding display label of the events selected from the given set of attributes to be presented with its attribute values on the y-axis of the cross plot, providing a pattern algorithm to detect whether an arrived event is part of a given pattern on the basis of a comparison of the attributes allocated to the given pattern and of the attributes assigned to the arrived event, providing a mapping algorithm to map any attribute value of an attribute selected from the given set of attributes onto the y-axis of the cross plot, allocating a second display label to the events indicating the attribute value of the attributes being uncovered as part of the given pattern, plotting all the events arrived within the time period and including an attribute value allocated to a primary attribute into the cross plot with the first display label indicating the primary attribute, and plotting all the events arrived within the time period and being detected by the pattern algorithm as part of the given pattern into the cross plot with the second display label indicating the given pattern.

30

[Fig. 2]

Figure 1

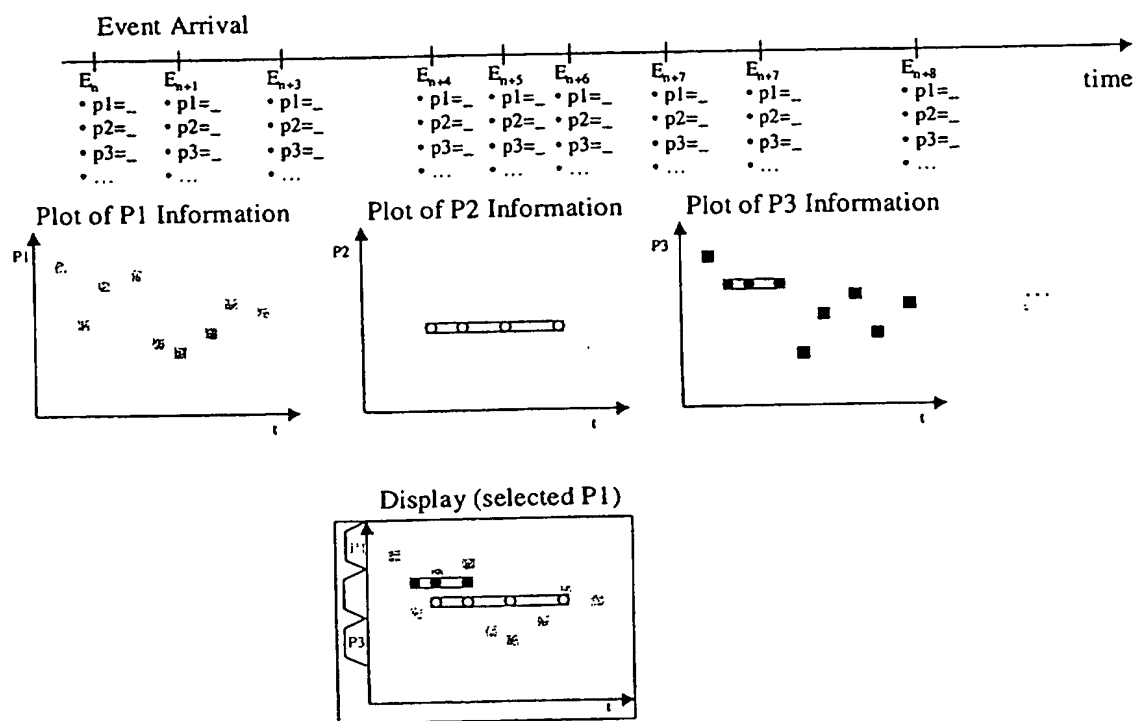


Figure 2

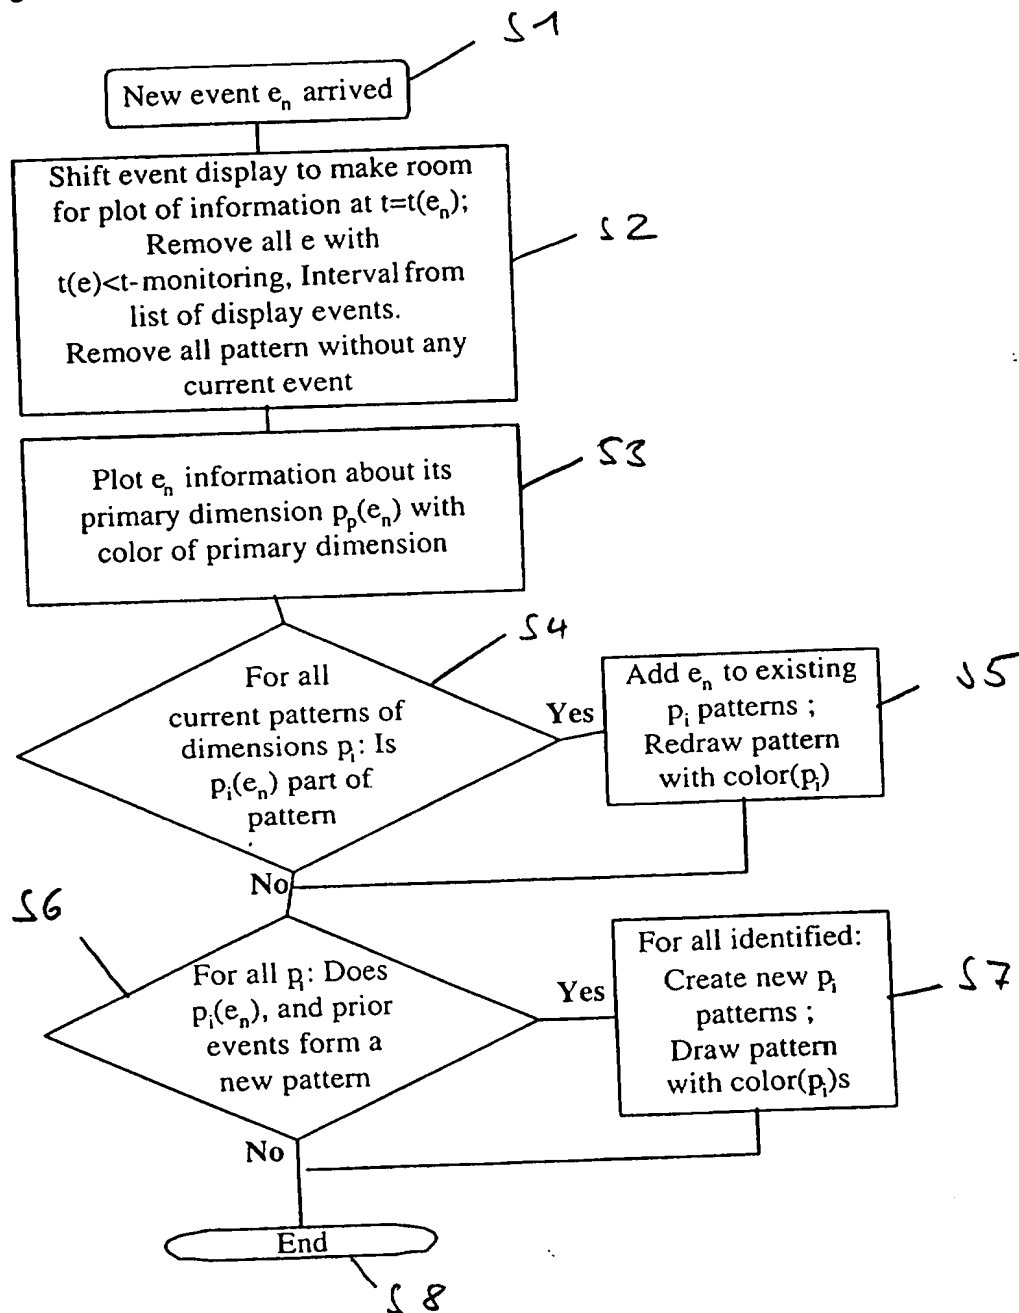


Figure 3

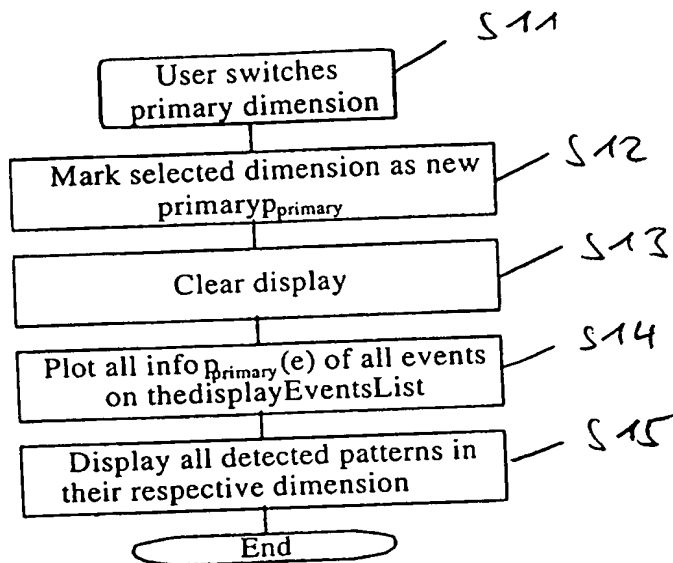


Figure 4

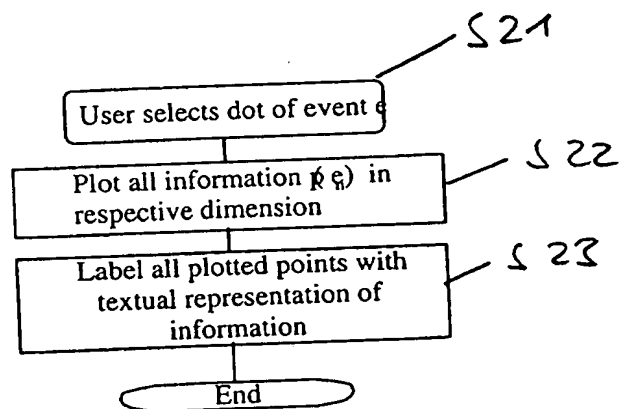


Figure 5

